

This deployment guide shows you how to deploy and connect to a Compute Engine virtual machine (VM) that is ready for the installation of SAP Adaptive Server Enterprise (ASE) on a Linux operating system.

The instructions in this guide use [Cloud Deployment Manager](/deployment-manager/docs/fundamentals) (/deployment-manager/docs/fundamentals) to deploy and configure the VM, the Linux operating system, and the required disk volumes for SAP ASE.

For details on planning your deployment, see the [SAP ASE Planning Guide](/solutions/sap/docs/sap-ase-planning-guide) (/solutions/sap/docs/sap-ase-planning-guide).

If you do not already have a Google Cloud project with billing enabled, you must create one before you can deploy a VM for your SAP ASE installation.

To create a project:

1. [Sign in](https://accounts.google.com/Login) (https://accounts.google.com/Login) to your Google Account.

If you don't already have one, [sign up for a new account](https://accounts.google.com/SignUp) (https://accounts.google.com/SignUp).

2. In the Cloud Console, on the project selector page, select or create a Cloud project.

★ **Note:** If you don't plan to keep the resources that you create in this procedure, create a project instead of selecting an existing project. After you finish these steps, you can delete the project, removing all resources associated with the project.

[Go to the project selector page](https://console.cloud.google.com/projectselector2/home/dashboard) (https://console.cloud.google.com/projectselector2/home/dashboard)

3. Make sure that billing is enabled for your Google Cloud project. [Learn how to confirm billing is enabled for your project](/billing/docs/how-to/modify-project) (/billing/docs/how-to/modify-project).

These instructions use [Cloud Shell](/shell/docs/) (/shell/docs/) to enter `gcloud` commands that deploy or configure your Google Cloud resources. Cloud Shell is accessed through the Cloud Console in your browser.

Cloud Shell runs on a VM that Google Cloud provisions each time you start Cloud Shell. The first time you use Cloud Shell, Google Cloud also creates a persistent `$HOME` directory for you, which is restored each time you open Cloud Shell.

The provisioned VM includes the latest [Cloud SDK](/sdk/docs/) (/sdk/docs/), which provides the `gcloud` command-line interface. Therefore, the `gcloud` commands that you use in Cloud Shell are the same as those you would use in a locally installed instance of the Cloud SDK.

If you have the Cloud SDK installed, you can issue the `gcloud` commands that are used in these instructions from your local machine. However, with a locally installed Cloud SDK you must always make sure that you are using the latest version of the Cloud SDK.

Whether you use Cloud Shell or Cloud SDK, you can set and change the properties of your `gcloud` command environment and save them as a *configuration*. Configurations are collections of key-value pairs that influence the behavior of the `gcloud` commands.

Some basic actions you can take with a configuration in Cloud Shell include:

- Initialize a configuration with `gcloud init`.
- Check the settings of your current `gcloud` configuration with `gcloud config list`.
- Change the Google Cloud project you are working in with `gcloud config set project [PROJECT_ID]` where `[PROJECT_ID]` represents your Google Cloud project.
- Set a default region with `gcloud config set compute/region [REGION]` where `[REGION]` represents a Google Cloud region.
- Set a default zone with `gcloud config set compute/zone [ZONE]` where `[ZONE]` represents a Google Cloud zone.
- Create a new configuration with `gcloud config configurations create [NAME]` where `[NAME]` represents the name for the configuration.

For more information about working with configurations, see [Managing Cloud SDK configurations](/sdk/docs/configurations/) (/sdk/docs/configurations/).

For security purposes, create a new network. You can control who has access by adding firewall rules or by using another access control method.

If your project has a default VPC network, don't use it. Instead, create your own VPC network so that the only firewall rules in effect are those that you create explicitly.

During deployment, VM instances typically require access to the internet to download Google's monitoring agent. If you are using one of the SAP-certified Linux images that are available from Google Cloud, the VM instance also requires access to the internet in order to register the license and to access OS vendor repositories. A configuration with a NAT gateway and with VM network tags supports this access, even if the target VMs do not have external IPs.

To set up networking:

1. Go to Cloud Shell.

[Go to Cloud Shell](https://console.cloud.google.com/?cloudshell=true) (<https://console.cloud.google.com/?cloudshell=true>)

2. To create a new network in the custom subnetworks mode, run:

where `[YOUR_NETWORK_NAME]` is the name of the new network. The network name can contain only lowercase characters, digits, and the dash character (-).

Specify `--subnet-mode custom` to avoid using the default auto mode, which automatically creates a subnet in each Compute Engine region. For more information, see [Subnet creation mode](/vpc/docs/vpc#subnet-ranges) (</vpc/docs/vpc#subnet-ranges>).

3. Create a subnetwork, and specify the region and IP range:

where:

- `[YOUR_SUBNETWORK_NAME]` is the new subnetwork.
- `[YOUR_NETWORK_NAME]` is the name of the network you created in the previous step.
- `[REGION]` is the region where you want the subnetwork. If you are using SAP ASE with SAP NetWeaver, use a region that is [supported for SAP NetWeaver](#)

(/solutions/sap/docs/netweaver-planning-guide#planning\_regions\_zones).

- **[YOUR\_RANGE]** is the IP address range, specified in CIDR format ([https://wikipedia.org/wiki/Classless\\_Inter-Domain\\_Routing](https://wikipedia.org/wiki/Classless_Inter-Domain_Routing)), such as 10.1.0.0/24. If you plan to add more than one subnetwork, assign non-overlapping CIDR IP ranges for each subnetwork in the network. Note that each subnetwork and its internal IP ranges are mapped to a single region.

4. Optionally, repeat the previous step and add additional subnetworks.

If you intend to create one or more VMs that will not have public IP addresses, you must create a NAT gateway so that your VMs can access the Internet to download Google's monitoring agent.

If you intend to assign an external public IP address to your VM, you can skip this step.

**Important:** Do not remove the public IP addresses from your new VMs until after the installation of your SAP software is complete and validated.

To create a NAT gateway:

1. Create a VM to act as the NAT gateway in the subnet you just created:

where:

- **[YOUR\_VM\_NAME]** is the name of the VM you are creating that want to use for the NAT gateway.
- **[YOUR\_ZONE]** is the zone where you want the VM.
- **[YOUR\_IMAGE\_FAMILY]** and **[YOUR\_IMAGE\_PROJECT]** specify the image you want to use (</compute/docs/images#os-compute-support>) for the NAT gateway.

- [YOUR\_MACHINE\_TYPE] is any supported machine type. If you expect high network traffic, choose a machine type with that has at least eight virtual CPUs.
- [YOUR\_SUBNETWORK\_NAME] is the name of the subnetwork where you want the VM.
- [YOUR\_VM\_TAG] is a tag that is applied to the VM you are creating. If you use this VM as a bastion host, this tag is used to apply the related firewall rule only to this VM.

2. Create a route that is tagged so that traffic passes through the NAT VM instead of the default Internet gateway:

where:

- [YOUR\_ROUTE\_NAME] is the name of the route you are creating.
- [YOUR\_NETWORK\_NAME] is the network you created.
- [YOUR\_VM\_NAME] is the VM you are using for your NAT gateway.
- [YOUR\_ZONE] is the zone where the VM is located.
- [YOUR\_TAG\_NAME] is the tag on the route that directs traffic through the NAT VM.

3. If you also want to use the NAT gateway VM as a bastion host, run the following command. This command creates a firewall rule that allows inbound SSH access to this instance from the Internet:

where:

- [YOUR\_NETWORK\_NAME] is the network you created.
- [YOUR\_VM\_TAG] is the tag you specified when you created the NAT gateway VM. This tag is used so this firewall rule applies only to the VM that hosts the NAT gateway, and not to all VMs in the network.

By default, an *implied firewall rule* blocks incoming connections from outside your Virtual Private Cloud (VPC) network. To allow incoming connections, set up a firewall rule for your VM. After an incoming connection is established with a VM, traffic is permitted in both directions over that connection.

You can create a firewall rule to allow external access to specified ports, or to restrict access between VMs on the same network. If the `default` VPC network type is used, some additional default rules also apply, such as the `default-allow-internal` rule, which allows connectivity between VMs on the same network on all ports.

Depending on the IT policy that is applicable to your environment, you might need to isolate or otherwise restrict connectivity to your database host, which you can do by creating firewall rules.

Depending on your scenario, you can create firewall rules to allow access for:

- ASE network listener for client connections. The default ASE port number is 5000.
- The ports required by SAP ASE Cockpit. For a list of ports that ASE Cockpit uses, see the [SAP ASE](https://help.sap.com/viewer/p/SAP_ASE) ([https://help.sap.com/viewer/p/SAP\\_ASE](https://help.sap.com/viewer/p/SAP_ASE)) Cockpit documentation in the SAP Help Portal.
- The default SAP ports that are listed in [TCP/IP of All SAP Products](https://help.sap.com/viewer/575a9f0e56f34c6e8138439eefc32b16/2.0/en-US/616a3c0b1cc748238de9c0341b15c63c.html) (<https://help.sap.com/viewer/575a9f0e56f34c6e8138439eefc32b16/2.0/en-US/616a3c0b1cc748238de9c0341b15c63c.html>)
- Connections from your computer or your corporate network environment to your Compute Engine VM instance. If you are unsure of what IP address to use, talk to your company's network administrator.
- Communication between VMs when, for example, your database server and application server are running on different VMs. To enable communication between VMs, you must create a firewall rule to allow traffic that originates from the subnetwork.
- SSH connections to your VM instance, including [SSH from the browser](https://cloud.google.com/compute/docs/ssh-in-browser) (<https://cloud.google.com/compute/docs/ssh-in-browser>).
- Connection to your VM by using a third-party tool in Linux. Create a rule to allow access for the tool through your firewall.

The following procedure is a simplified version of the instructions for creating firewall rules. For more detailed instructions, see [Virtual Private Cloud documentation](#) ([/vpc/docs/using-firewalls#creating\\_firewall\\_rules](/vpc/docs/using-firewalls#creating_firewall_rules)).

To create a firewall rule:

1. In the Cloud Console, go to the **Firewall rules** page.

**OPEN FIREWALL RULES** (<https://console.cloud.google.com/networking/firewalls/list>)

2. At the top of the page, click **Create firewall rule**.

- In the **Network** field, select the network where your VM is located.
- In the **Targets** field, specify the resources on Google Cloud that this rule applies to. For example, specify **All instances in the network**. Or to limit the rule to specific instances on Google Cloud, enter tags in **Specified target tags**.
- In the **Source filter** field, select one of the following:
  - **IP ranges** to allow incoming traffic from specific IP addresses. Specify the range of IP addresses in the **Source IP ranges** field.
  - **Subnets** to allow incoming traffic from a particular subnetwork. Specify the subnetwork name in the following **Subnets** field. You can use this option to allow access between the VMs in a 3-tier or scaleout configuration.
- In the **Protocols and ports** section, select **Specified protocols and ports** and enter **tcp: [PORT\_NUMBER]**.

3. Click **Create** to create your firewall rule.

The following instructions use Deployment Manager to deploy a VM instance with Linux and all of the persistent disks that SAP ASE requires. You define the values for the installation in a Deployment Manager configuration file template.

Deployment Manager treats all of the resources that are created for the SAP ASE system as a single entity called a *deployment*. You can view all of the deployments for your Google Cloud project on the [Deployments](https://console.cloud.google.com/dm/deployments) (<https://console.cloud.google.com/dm/deployments>) page in the Cloud Console.

Deleting a deployment deletes all of the resources associated with the deployment, including the SAP ASE system, the persistent disks that were deployed with it.

The following instructions use Cloud Shell, but if you prefer, you can use the Cloud SDK in your local terminal instead.

1. Open the Cloud Shell.

[Go to the Cloud Shell](https://console.cloud.google.com/?cloudshell=true) (https://console.cloud.google.com/?cloudshell=true)

2. Download the template.yaml configuration file template to your working directory by entering the following command in the Cloud Shell:

3. Optionally, rename the template.yaml file to identify the configuration it defines.

4. Open the template.yaml file in the Cloud Shell code editor.

To open the Cloud Shell code editor, click the pencil icon in the upper right corner of the Cloud Shell terminal window.

5. In the template.yaml file, update the following property values by replacing the brackets and their contents with the values for your installation.

Property	Data type	Description
instanceName	String	The name of the VM instance on which SAP ASE will be installed. The name must be 13 characters or less and be specified in lowercase letters, numbers, or hyphens.
instanceType	String	The type of Compute Engine <a href="#">virtual machine</a> (/compute/docs/machine-types) on which SAP ASE will be installed. If you are running SAP ASE with SAP NetWeaver on the same VM, select a machine type that includes enough CPUs and memory to support both systems. See the <a href="#">SAP NetWeaver Planning Guide</a> (/solutions/sap/docs/netweaver-planning-guide#machine_types).
zone	String	The zone in which you are deploying your SAP ASE. It must be in the same region that you selected for your subnetwork.
subnetwork	String	The name of the subnetwork that you created in a previous step. If you are deploying to a shared VPC, specify this value as <b>[SHAREDVPC_PROJECT]/[SUBNETWORK]</b> . For example, myproject/network1.



Property	Data type	Description
linuxImage	String	The name of the Linux operating- system image or image family that you are using with SAP ASE. To specify an image family, add the prefix <code>family/</code> to the family name. For example, <code>family/rhel-7</code> or <code>family/sles-12-sp2-sap</code> . To use a specific image, enter only the image name. For the list of available image families, see the <a href="https://console.cloud.google.com/compute/images">Images</a> ( <a href="https://console.cloud.google.com/compute/images">https://console.cloud.google.com/compute/images</a> ) page in the Cloud console.
linuxImageProjectString		The Google Cloud project that contains the image you are going to use. This project might be your own project or a Google Cloud image project, such as <code>rhel-sap-cloud</code> or <code>suse-sap-cloud</code> . For a list of GCP image projects, see the <a href="/compute/docs/images">Images</a> ( <a href="/compute/docs/images">/compute/docs/images</a> ) page in the Compute Engine documentation.
aseSID	String	The database instance ID.
asesidSize	Integer	The size in GB of <code>/sybase/[DBSID]</code> , which is the root directory of the database instance. In the deployed VM, this volume is labeled <b>ASE</b> .
aseddiagSize	Integer	The size of <code>/sybase/[DBSID]/sapdiag</code> , which holds the diagnostic tablespace for SAPTOOLS.
asesaptempSize	Integer	The size of <code>/sybase/[DBSID]/saptmp</code> , which holds the database temporary table space.
asesapdataSize	Integer	The size of <code>/sybase/[DBSID]/sapdata</code> , which holds the database data files.
aseLogSize	Integer	The size of <code>/sybase/[DBSID]logdir</code> , which holds the database transaction logs.
asebackupSize	Integer	The size of the <code>/sybasebackup</code> volume. If set to 0 or omitted, no disk is created.
asesapdataSSD	boolean	The SSD toggle for the data drive. If set to true, the data disk will be SSD.
aseLogSSD	boolean	The SSD toggle for the log drive. If set to true, the log disk will be SSD.
usersapSize	Integer	Required only if you are installing SAP ASE to run with SAP NetWeaver on the same VM instance.
sapmntSize	Integer	Required only if you are installing SAP ASE to run with SAP NetWeaver on the same VM instance.
swapSize	Integer	Required only if you are installing SAP ASE to run with SAP NetWeaver on the same VM instance.

The following configuration file creates a VM that is configured to run both the SAP ASE database server and SAP NetWeaver. The configuration file directs Deployment Manager to deploy an `n1-standard-16` VM that is running a SLES 12 SP2 operating system. The VM includes all of the directories that are required to run SAP ASE with NetWeaver.

## 6. Create the VM instance:





The above command invokes the Deployment Manager, which deploys the VM according to the specifications in your template.yaml file. The process might take a few minutes. To check the progress of your deployment, follow the steps in the next section.

1. Open Stackdriver Logging to check for errors and monitor the progress of the installation.

★ **Note:** You might incur costs when completing this step in Stackdriver. For more information, see [Stackdriver pricing \(/stackdriver/pricing\\_v2\)](#).

[Go to Logging](https://console.cloud.google.com/logs/viewer) (https://console.cloud.google.com/logs/viewer)

- On the Resources tab, select **Global** as your logging resource. If "**INSTANCE DEPLOYMENT COMPLETE**" is displayed for a VM, Deployment Manager processing is complete for the VM.

▶		2018-05-21 15:27:19.752 PDT	ex-vm-ase-lin	Deployment	"--- /sybasebackup"
▶		2018-05-21 15:27:36.615 PDT	ex-vm-ase-lin	Deployment	"Creating file systems for SAP NetWeaver"
▶		2018-05-21 15:27:37.336 PDT	ex-vm-ase-lin	Deployment	"--- /usr/sap"
▶		2018-05-21 15:27:41.484 PDT	ex-vm-ase-lin	Deployment	"--- /sapmnt"
▶		2018-05-21 15:27:45.626 PDT	ex-vm-ase-lin	Deployment	"--- swap volume"
▶		2018-05-21 15:27:46.459 PDT	ex-vm-ase-lin	Deployment	"INSTANCE DEPLOYMENT COMPLETE"

- After the SAP ASE system deploys without errors, connect to your VM by using SSH. From the Compute Engine [VM instances page](https://console.cloud.google.com/compute/instances) (https://console.cloud.google.com/compute/instances), you can click the SSH button for your VM instance, or you can use your preferred SSH method.

<input type="checkbox"/>	Name ^	Zone	Internal IP	External IP	Connect
<input checked="" type="checkbox"/>	ex-vm-ase-lin	us-central1-f	10.1.0.4 (nic0)	104.198.240.61	SSH

- Change to the root user.

- At the command prompt, enter `df -h`. Ensure that you see output similar to the following, such as the `//sybase/[DBSID]/sapdata` directory.

```
ex-vm-ase-lin:~ # df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  30G         0    30G   0% /dev
tmpfs                     45G         0    45G   0% /dev/shm
tmpfs                     30G        18M    30G   1% /run
tmpfs                    30G         0    30G   0% /sys/fs/cgroup
/dev/sda3                 30G        2.3G    26G   9% /
/dev/sda2                 200M        660K   200M   1% /boot/efi
/dev/mapper/vg_asesid-vol  9.8G        23M    9.2G   1% /sybase/AS1
/dev/mapper/vg_asesapdata-vol 50G        52M    47G   1% /sybase/AS1/sapdata_1
/dev/mapper/vg_asesaplog-vol 30G        45M    28G   1% /sybase/AS1/saplog_1
/dev/mapper/vg_asesaptemp-vol 20G        45M    19G   1% /sybase/AS1/saptemp
/dev/mapper/vg_asesapdiag-vol 7.8G       19M    7.4G   1% /sybase/AS1/sapdiag
/dev/mapper/vg_asebackup-vol 99G        60M    94G   1% /sybasebackup
/dev/mapper/vg_usrsap-vol  15G        38M    14G   1% /usr/sap
/dev/mapper/vg_sapmnt-vol  15G        38M    14G   1% /sapmnt
```

- Confirm that the swap directory was created by entering the following command:

You should see results similar to the following example:

```
ex-vm-ase-lin:~ # cat /proc/meminfo | grep Swap
SwapCached:          0 kB
SwapTotal:           25161724 kB
SwapFree:            25161724 kB
```

If any of the validation steps show that the installation failed:

1. Correct the error.
2. On the [Deployments](https://console.cloud.google.com/dm/deployments) (https://console.cloud.google.com/dm/deployments) page, delete the deployment to clean up the VMs and persistent disks from the failed installation.
3. Rerun your deployment.

Now that your operating system is configured, you can install your SAP ASE database:

To install SAP ASE on your VM:

1. Establish an SSH connection to your Linux-based VM.
2. Download or copy your SAP ASE installation media from SAP to your VM.
3. Prepare your SAP ASE SySAM license file or license server for use with the installation.
4. Install the SAP ASE database. See the [SAP ASE Installation and Upgrade Guide for Linux](https://help.sap.com/viewer/23c3bb4a29be443ea887fa10871a30f8/latest) (https://help.sap.com/viewer/23c3bb4a29be443ea887fa10871a30f8/latest) for guidance.

Before using your SAP ASE instance, we recommend that you perform the following post-deployment steps:

1. Update your SAP ASE software with the latest patches, if available.
2. Install any additional components.

### 3. Configure and back up your new SAP ASE database.

For additional post-deployment guidance, see the Post-installation Tasks section of the [SAP ASE Installation and Upgrade Guide for Linux](#)

(<https://help.sap.com/viewer/23c3bb4a29be443ea887fa10871a30f8/latest/en-US/a6e8dcfbbc2b1014a59bceaf9d7f47ce.html>)

This section contains information about how to correct common issues.

If you are having issues connecting to your VM through SSH, ensure that you have created a firewall rule to open port 22 on the GCP network you are using.

For other possible issues, see [Known issues](#) (/compute/docs/ssh-in-browser#known\_issues) for SSH from the browser.